Challenges of Embodied Intelligence

(Invited Paper)

Janusz A. Starzyk, Yinyin Liu, and Haibo He
School of Electrical Engineering and Computer Science
Ohio University
Stocker Center, Athens, OH, 45701 USA
starzyk@bobcat.ent.ohiou.edu

***Abstract*—This paper focuses on challenges of designing working models of embodied intelligence (EI). Formal definitions of embodied intelligence and its embodiment are given. Three self-organizing hierarchical structures – sensory, motor, and goal creation pathways form the core of EI. The three pathways are developed simultaneously extracting knowledge and learning skills through interaction with the environment. Technological challenges of building human-level intelligence are discussed.**

I. Embodied Intelligence

There is no agreement how to define intelligence; however, there is a good understanding of what an intelligent agent (biological, mechanical or virtual) must be capable of. Scientists list such capabilities as abstract thinking, problem solving, intuition, creativity, consciousness, emotions, learning, memory and motor skills as traits of intelligence. They use various tests and measures to compare the levels of intelligence and differentiate between intelligence of humans and other species. In fact, passing various tests for (human level) intelligence was used as a substitute for its definition. Complex skills and behaviors were used to define how intelligence manifests itself. This was a result of poor understanding of what is needed to create intelligence. Such approach was inconsistent, because once a machine that was obviously not intelligent satisfied one test, another test was used in its place.

In order to build working models of intelligent machines, an arbitrary and utilitarian definition of intelligence is adopted in this work. We will demonstrate that the definition is general enough to characterize agents of various levels of intelligence including human. To differentiate it from enigmatic meaning of intelligence, we will limit it to embodied intelligence suggested by Brooks [1] and described in more detail by Pfeifer [2].

Embodied intelligence (EI) has developed into a multidisciplinary field, including biology, neuroscience, electrical engineering, robotics, biomechanics, material science, and dynamic systems. It focuses on understanding biological intelligent systems, extracting general principles of intelligent behavior and applying this knowledge to design robots and intelligent devices. Traditional artificial intelligence (AI) has focused on developing computational aspects of intelligence, looking at cognition as a form of computation, and developing special high level skills (like logical reasoning or theorem proving) in abstraction of embodiment and environment. However, all natural intelligent systems have biological bodies and are situated in a set environment. The working principles of a biological system can not be fully understood in purely computational terms as they are functions of its environment, its ecological niche, and its evolutionary history. This includes the type and the number of sensors it uses, the actuators it employs, and the dynamics of its physical body.

It is our aim to base the design concepts of embodied intelligence on a minimum set of requirements and mechanisms from which all traits of intelligence can be derived.

Definition:

Embodied intelligence (EI) is defined as a mechanism that learns how to survive in a hostile environment.

A mechanism in this definition applies to all forms of embodied intelligence, including biological, mechanical or virtual agents with fixed or variable embodiment, fixed or variable sensors and actuators. Implied in this definition is that EI interacts with environment and that the results of actions are perceived by its sensors. Also implied is that the environment is hostile to EI so that EI has to learn how to survive. This hostility of environment symbolizes all forms of pains that EI may suffer – whether it is an act of open hostility or simply scarcity of resources needed for the survival of EI. The important fact is that the hostility is persistent. For example, battery power is a persistent threat for an agent requiring it. Gradually the energy level goes down, and unless the EI replenish its energy, a perceived discomfort from the its energy level sensor will increase. This perpetual hostility will be a foundation for creating a value system and the goals that the EI sets for itself while interacting with its environment.

A critical element of EI is learning. Thus an agent that knows how to survive in a hostile environment but cannot learn new skills is not intelligent. Learning to survive requires not only memory but its management, so that only the important memories must be retained. Learning also requires ability to associate sensory and motor signals so that outcomes can be linked with causes.

Hostile action of the environment towards EI is necessary for it to develop environment related skills, build models of the environment and its embodiment, explore and learn successful actions. Thus, to paraphrase a quote by the Field Crown Hetman, Stefan Czarniecki "I am not from the reaches of the land but from what pains me", (in Polish – jam nie z soli ani z roli jeno z tego co mnie boli), pain is necessary to develop intelligence. In our model, pain and pain management in EI will become a foundation for building a value system, learning, goal creation, planning, thinking, problem solving, creativity, and developing motor skills. In more advanced forms of EI it will also lead to intuition, consciousness, and emotions. Thus all forms and levels of intelligence can be considered under the proposed definition of EI.

Notice that this definition of EI clearly differentiates knowledge from intelligence. While knowledge is the acquired set of skills and information about the environment, intelligence requires the ability to acquire knowledge.

In this work we will show how some elements of EI can be implemented in a hierarchical, self-organizing, learning network of processing elements called neurons. We hope that this paper will contribute to building working models of EI.

II. EMBODIMENT AND INTELLIGENCE

Intelligence cannot develop without an embodiment. The intelligence core interacts with the environment through its embodiment, as shown in Fig. 1.

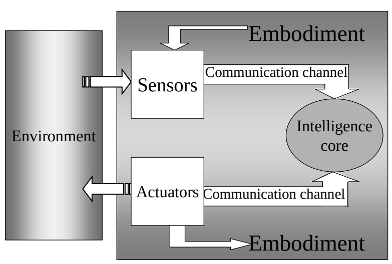

Fig. 1. Intelligence core with its embodiment and environment

The embodiment is at the same time a part of the environment that can be perceived, modeled and learned by intelligence, as it is an extension of the EI interface that interacts with the environment. Through this extension of its interface, which could be in the form of motor actions, the EI affects the environment. The response of the environment is registered through sensors implanted in the embodiment. Properties of these sensors, their status and limitations can be studied, modeled and understood by EI.

However, the embodiment does not have to be constant nor physically attached to a body. The boundaries between embodiment and the environment change during the interaction which modifies the EI's self-determination. This interaction can be viewed as the closed-loop sensory-motor coordination, which will be discussed in detail in later sections. Because the boundaries between embodiment and environment are dynamically changing, the definition of embodiment has to reflect this fact and contain elements of indetermination.

Definition:

Embodiment of EI is a mechanism under the control of the intelligence core that contains sensors and actuators connected to the core through communication channels.

A first consequence of this definition is that the mechanism under control may change. For instance, if one loses an arm its embodiment changes in an obvious way. The embodiment may also change due to a disease, affecting the way it works or how it interacts with the environment, etc.

Secondly, embodiment does not have to be permanently attached to intelligence in order to play its role of sensory-motor interaction with the rest of the environment. For instance, if we operate a machine (drive a car, use keyboard, play tennis), our embodiment dynamics can be learned and associated with our action to the extent that reduces the distinction between the dynamics of our own body and the dynamics of our body operating in tandem with the machine. Likewise, artificially enhanced senses can be perceived and characterized as our own senses (e.g. glasses that improve our vision or a hearing aid that improve our hearing). Another example of sensor extension could be an electronic implant stimulating the brain of a blind person to provide visual information.

Extended embodiment does not have to be of a physical (mechanical) nature. It could be in the form of remote control of tools in a distant surgery procedure or monitoring Martian landscape through mobile cameras. It could also be our distant presence at the soccer game through received TV images or our voice message delivered through a speakerphone to a group of people at the teleconference.

Finally, extended embodiment of intelligence comes in the form of organizations and their internal working mechanisms and procedures. A general directing troops on a battle field feels a similar power of moving armies as a crane operator feels the mechanical power of the machine while he operates. The president also feels the power of his address to the nation and the large impact it makes on people's lives.

This extended embodiment enhances EI's ability to interact with the environment and thus its ability to grow in complexity, skills and effectiveness. If the President learns how to address the nation, his ability to affect the environment grows differently than that of a woman in Darfur trying to save her child from violence.

Our knowledge of embodiment properties is a key to its proper use in interaction with the world. We rely on this knowledge to plan our actions and predict the responses from the environment. A change in the way our embodiment implements desired actions or perceives response from the environment introduces uncertainty in our behavior and may lead to confusion and less than optimum decision making. If a car's controls were suddenly reversed during operation a user would require some adaptation time to adjust to the new situation and probably would not be able to before crashing. Therefore, what we learn about our environment and our ability to change this environment is affected not only by our intelligence (ability to learn, understand, represent, analyze and plan) but by correct perception of our embodiment as well.

III. Designing the Embodied Intelligence

Learning is an active process. EI acquires information about its environment through sensors and interacts with it by sensory-motor coordination. The motor neurons fire in response to excitations provided by sensory and goal creation neurons according to desired actions associated with the perceived situation. The agent learns which actions are desirable and those that are not by using a value system. One mechanism to build an internal value system uses the reinforcement learning [3].

Learning which actions are desirable and which are not makes the learning agent more fit to survive in the hostile environment. There are several means of adapting to the environment that an agent can use to survive: evolutionary - by using the natural selection of those agents that are most fit; physiological - by developing new motor skills like sweating in the hot weather or new sensors like cell sensitivity to light; and cognitive - by learning, using spatio-temporal memory, pattern recognition, and associations. Here we address only the later, and the most critical, form of adaptation for the development of EI.

R. Brooks [1][4], the father of embodied intelligence, proposed to design this system through layers of simple sensory-motor coordination built on finite state machines (FSM). In his subsumption architecture higher levels are built upon the lower levels subsuming the lower levels functionality. In subsumption architecture, each layer consists of asynchronous modules that send messages to each other. Each module is an augmented FSM. Inputs to such modules can be suppressed and outputs inhibited by signals from other modules.

Although a subsumption architecture may be an efficient design approach to build robots with complex behavior, it cannot lead to intelligence. A designer must be involved in developing each FSM. These FSMs do not know how to modify their own structures to handle new tasks. There is no self-organization and no learning.

Moreover, since new tasks may not be compatible with old ones, modification of the machine behavior to incorporate new tasks may become extremely difficult. Very quickly complexity exceeds understanding of the machine's operation by the human designer, who no longer understands how to add a new layer of functionality.

R. Pfeifer [2] modified the subsumption architecture approach to include self-organization and the emergence of necessary links between lower level processes that control sensory-motor coordination. He also added a value principle to his design approach, requiring a mechanism for self-supervised, perpetual learning that employs the principle of self-organization. The value system acts as a teacher telling an agent what actions are good for its objectives. Memory of the recent history is necessary to implement this value system. This memory is accomplished by time-averaging neuron activities. But as Sporns and Edelman pointed out [5] "the issue of value constraints and their number presents one of the greatest future challenges to selectional theories of brain function." In our work we try to address this issue.

A practical effort to design structural and algorithmic properties of the neocortex was undertaken by J. Hawkins from Numenta Inc.[6]. Numenta develops software code for Hierarchical Temporal Memory (HTM). HTM uses a hierarchy of spatio-temporal associations and learns complex goal oriented behaviors. The information, in the form of probability distributions, passes up and down the hierarchy to represent the sensory inputs and to make predictions. It uses a combination of unsupervised and supervised learning to make associations. In the authors' opinion HTM may yield machines that exceed human level performance in cognitive tasks.

Earlier attempts to design working models of intelligence include for instance GOMS [7], SOAR [8][9], and ACT-R [10] software systems. GOMS (Goals, Operators, Methods, and Selection Rules) is a software system for modeling and describing human performance that provides a framework to analyze human computer interactions. Goals, that a user is trying to accomplish, are organized hierarchically. Methods describe sequences of basic operations used to accomplish a goal. Selection rules describe which method should be used in a particular situation. It uses a production-system representation of human procedural knowledge required to accomplish production goals. It gives good quantitative predictions of performance time and learning.

SOAR (State, Operator, Application, Result) is a cognitive goal-oriented architecture that develops a minimal set of rules to support intelligent behavior in a specified environment. It uses symbolic knowledge and knowledge-based symbolic reasoning to solve problems. It creates subgoals even with incomplete or inconsistent knowledge. SOAR can also generate rules for the implementation of goals using a process called chunking. The SOAR program learns using explanation-based learning, macro-operator learning, strategy acquisition, and learning by instruction.

ACT-R (Adaptive Control of Thought - Rational) is a model of the human cognitive process focusing on learning and problem solving. Cognitive tasks are performed using if-then production rules, with working memory (declarative or procedural). ACTR-R uses pattern matching to match conditions for its production rules and conflict resolution to decide which rule applies. Using ACT-R requires developing a domain-specific knowledge model of the cognitive task for a specific application. None of these three systems uses self-organization, unsupervised learning, or creates a knowledge base for its actions with the environment.

A. Basic requirements for EI

The way that the brain stores a pattern in its hierarchical memory of the neocortex is very different from the way a computer does. Neurons in the human brain self-organize to store the patterns to which the brain is exposed. Human utilizes his senses to build a perception of the environment and activates appropriate motor neurons to apply actions. This enables a human to build a model of the world in a fascinating way. He uses this model to quickly recognize patterns in order to respond to the external stimuli and interact within the environment. He also uses this model to expect future events, accomplish efficient planning, and do logical thinking. The perceptions and actions are activated selectively by the brain with attention focused on those observations and actions that are related to human objectives.

Thus a goal driven behavior is one of the required elements of intelligence. In addition, since humans and animals create their own goals, it is desirable for the EI to be able to do so too. While the goal creation mechanism in a human is not obvious to behavioral scientists or psychologists, it may be one of the most important elements of EI mechanism.

In the existing models for designing of intelligent agents, the goal is defined by designers and is given to the learning agent. It is desired that the agent finds a way to achieve the goal by its own actions. Having domain specific knowledge, an agent may be allowed to formulate subgoals to achieve goals as in the SOAR architecture. During the process of finding solutions, the agent will build a value system which evaluates different available actions and chooses the best one according to their values. In such cases, the agent is not able to create its own goals or find the sub-goals in order to accomplish a complicated task. In fact, we would argue that an agent who only follows externally set goals would not be able to develop some higher level cognitive abilities such as deliberate thought, free will, intuition, consciousness or emotions. To some degree it will be more similar to a robot than to an enlightened individual. Therefore, we set goal creation ability as one of the fundamental requirements for EI.

Setting a goal makes the machine specialize to perform in specific types of operation. While this may be a useful limitation from a utilitarian point of view, and the resulting machine may be more efficient in implementing a set task, such a limitation will make the machine less intelligent. Such a solution would be sidestepping from the main goal of designing human level intelligence.

Since we require that an intelligent machine must have a built-in mechanism to create goals for its behavior, development of a goal creation system (GCS) should be one of our first priorities. We desire to build a GCS based on a simple and uniform structure interacting with its sensory and motor functions. In a sense, goal creation should result from the machine's interaction with its environment, by perceiving successes or failures of its actions.

As we desire to develop the machine with self-organizing, hierarchical structures and enable it to make spatiotemporal associations between sensory and motor functions, we would also expect the machine to use similar mechanisms to self-organize the goal creation.

We propose that in order to build intelligent machines, the following elements are essential:

1. **Hierarchical self-organizing learning memory (HSOLM)** to perceive and act according to machine's objectives.
2. **Goal creation system (GCS)** to develop sensory-motor coordination (SMC) and to act as stimuli for interaction with the environment.

EI must be able to self-organize its learning using internal mechanisms to acquire knowledge that is useful for its interaction with its environment. In the proposed model of EI, HSOLM will use three basic pathways – a **sensory pathway** responsible for perception, a **motor pathway** responsible for actions, and a goal creation pathway. The **goal creation pathway** is responsible for goal creation, evaluation of actions in relation to its goals, learning of useful associations and stimulating machine to perform useful actions. These three pathways interact on various abstraction levels, providing associations between the sensory, motor and goal creation pathways.

Goal creation stimulates the growth of hierarchical structures representing sensory inputs, actions and skills acquired by the machine, and abstract goals that the machine creates for itself. EI learns predominantly in an unsupervised manner by responding to stimuli from the environment. Learning is deliberate, perpetual, and related to satisfactory completion of EI goals. Hostile stimulation from the environment is necessary for the EI to grow in sophistication and to acquire necessary knowledge and skills.

B.Hierarchical self-organizing learning memory

The spatio-temporal patterns that a human experiences during his lifetime constitute the knowledge stored in his memory. The patterns have features represented on different abstraction levels, so they should be organized hierarchically based upon the abstraction levels.

Such patterns will be accordingly stored in the processing units (neurons and their connections) on different hierarchical levels of HSOLM. Thus, neurons on different levels handle the recognition tasks with different level of abstraction. Lower-level neurons are either activated directly by the sensory inputs or recognize certain detailed features. Subsequent level neurons combine the extracted features and represent elements of more complex entities. The information is gathered, associated and abstracted (in an invariant form) as it flows upwards in the hierarchy. Finally, top-level neurons represent perceived entities.

The human cognition is a very efficient process. Typically, it takes less than one second for sensory information to be obtained, perceived, processed and acted upon. Since neurons in the human brain take several milliseconds to fire, a typical cognition task takes less than 100 sequential pattern-processing steps. This recognition performed within such short period indicates the high efficiency of human perception. Such recognition along with the corresponding action is a very fundamental task the human can handle; nonetheless, it is already quite complicated for existing AI machines.

Mountcastle argues that all regions of the brain are built from the same structural units that perform the same computational algorithm [11]. These vertical column-like units called minicolumns (or microcolumns) contain groups of neurons connected locally in a pseudorandom way [12][13]. Their existence indicates system organization in clusters of neurons with possibly higher and more robust computational ability than a single neuron. Minicolumn organization helps to structure the network of neurons and increase sparsity of interconnections, which are useful in hardware implementation.

Another critical aspect of human brain development is self-organization. As the brain develops and learns it changes its own structure. A six year old child has many redundant and plastic connections ready to learn almost anything. After years of learning, the connection density among neurons is reduced, as only the most useful information is retained, and related memories and skills are refined. In a sense, neurons self-organize to store information. As a child turns from a novice to an expert, he also looses some ability to learn. To some degree learned knowledge obstructs having a fresh look at the old problem as is nicely expressed in the proverb “You can’t teach an old dog new tricks.”

Thus self-organization is a critical requirement for building structural elements of EI. It is important for three reasons. First, it is a local, learning driven process, thus it can be implemented in a massively parallel way through interacting neurons. The second reason is that by self-organization the machine can design its own structures without need for a detailed blueprint. Such detailed blueprint would be very difficult to design, build, test, or maintain. The opposite is true – we may be able to design a complex brain like structure with relatively little information encoded in its building blocks. The third reason is that it contains elements of error correction and self-repairing. If a single connection or neuron fails, its function will be replaced by other nearby neurons or by newly created connections.

The HSOLM is built as a self-organizing hierarchical structure with minicolumns as fundamental **processing units**. Different layers of the hierarchy represent different levels of abstraction. Inside the minicolumn, the neurons are interconnected in a predefined way. Minicolumns are connected with other minicolumns using connections both from different hierarchical levels and from the same level. Thus we may identify feedforward and feedback connections between different levels and associative connections on the same level. The architecture of such hierarchical memory is shown in Fig. 2 and contains a number of minicolumns on each level. The minicolumns contain a number of internal connections, associative connections to other minicolumns on the same level, feedforward connections from minicolumns on lower levels to those on higher levels and feedback connections from the higher levels to the lower levels. In such structures the three types of pathways (sensory, motor and goal creation) will be built and will interact with each other.

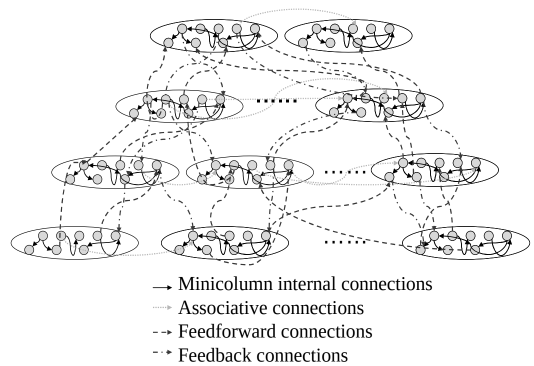

Fig. 2. Hierarchical structure to implement an HSOLM memory

In HSOLM, there are many more units on the higher level than on the lower levels in order to store a large number of entities. The feature recognized in the lower-level minicolumn may be related to multiple objects represented on the higher-level, so the lower-level neurons’ activity may be spread to several neurons on the higher level. However, the learning algorithm is set to restrict such activities. Therefore, we expect fewer active units on the higher levels than on the lower levels.

1)Sensory Pathway

The primary objective of neurons in the sensory pathway is to register the input information received from the environment and to build intentional representations to either be acted upon or stored in long-term memory for later use. The intentional representation is an internal representation of objects, symbols, abstractions, relations, actions, etc. (jointly called **entities**) related to EI’s interaction with the environment. Thus intentional representations have to be related to the machine’s objectives. EI focuses its attention on representation and recognition of entities that are important to its goals, values, and motor actions. These emerge gradually through machine’s operations and through its goal creation and value system. The machine determines which entities and related actions need to be represented in its memory. Intentional representations generalize the information received by introducing some degree of invariance which increases with the level of generalization. Representations are built using a self-organizing hierarchy of spatio-temporal patterns, that are a result of deliberate actions and learning.

Neurons in the sensory pathway are organized in a semi-hierarchical way to be able to register a large number of intentional representations. At the bottom of the hierarchy lie the raw sensory inputs. Neurons' activation on these inputs may represent a large number of input patterns that the machine receives over time. Each layer following the input layer both increases in overall size as well as has a reduced number of firing neurons. The result is a very large increase in storage capacity with each added layer.

Memories are developed through the modification of the interconnection weights between active neurons. Hierarchical organization also helps to streamline the input data representations, yielding large capacity memory structure and efficient signal processing scheme for received data. This system of hierarchical structuring of neurons is important for several reasons.

First, it provides the system with the ability to learn at very high capacities. Weight adjustment has an impotent impact on the machines ability to learn. Our system restricts the magnitude of weight adjustment based on the amount of times a neuron's weights have been previously modified. As a result of the expanding layers and reduced probability of firing, weights on the lower levels are adjusted more frequently then upper level weights. As sensory information is presented to the network, lower level connections quickly become rigid and can no longer be modified (learn). In contrast, upper level neurons are infrequently modified and therefore retain their ability to be modified even after extensive exposure to input level stimuli. This results in a system which can be trained to form representations on each increasing level with an increased ability to represent new objects.

Another reason why this hierarchical organization of the sensory pathway is important is related to power dissipation. Learning and signal processing consume energy. Power dissipation is one of the most critical design factors, especially in large, parallel computing systems (discussed in the section on challenges of EI implementation.) Thus, a learning and processing model that involves only a small subset of active neurons will save energy. A schematic representation of activations of the processing units in the sensory pathway is shown in Fig. 3.

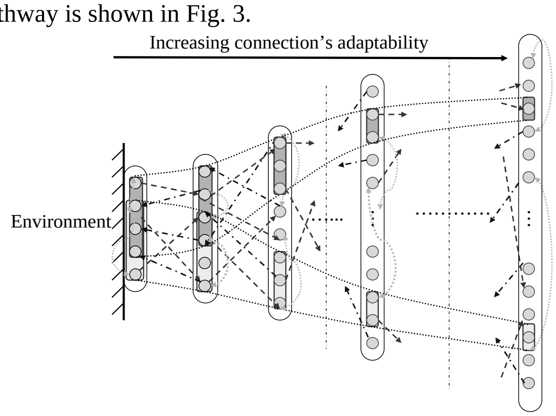

Fig. 3. Example activations of units in the sensory pathways

Two such activated pathways are shown to indicate reuse of the same lower level features in the recognition process. The figure also shows the increasing number of higher level units and the various types of connections among them.

EI uses two mechanisms to store the information in its memories and to build invariances. One is screening for novelty. Only new and useful information is stored in the form of **intentional representations**. To accomplish this, the machine continuously predicts what information will be coming in. This prediction manifests its understanding of the perceived signals and an assumption of sameness of the observed scene. Sensory input changes, are usually the result of minor changes in the point of view, motion, lighting conditions, or gradual modification of the object shape, color or form. When EI interacts with the environment, it is situated in a specific location, observing specific objects, and performing a specific task. Thus the assumed continuity of the observed sensory input is used for invariance building. Although we look at an object from different angles, we know that this is the same object and thus these various inputs must trigger the same representation on the higher level.

The built-in goal creation and value system provides a mechanism that triggers learning of intentional representations and associations between sensory and motor pathways. When the EI realizes that a specific action resulted in a desirable effect it stores the representation of the perceived entity and learns associations between the activated representations in the sensory pathway and the neurons in the motor pathway responsible for this action. If the effect is not desirable, it learns not to perform such an action. Finally, when no goal is affected, no learning takes place. The machine does not create intentional representation nor does it remember the action it took. Since this happens most of the time, such organization of the learning process protects the machine's memory from overloading of unimportant information. In addition, no learning is required if the obtained positive (or negative) outcome was expected. At most, some incremental changes in the interconnection strength may be observed in such cases.

2)Motor Pathway

The primary objective of neurons in the motor pathway is to represent and control execution of actions. The motor pathway represents skills learned by the EI. EI interacts with environment through actions, controlled by spatio-temporal sequences of motor neurons firing. The ability to perform these actions (skills) emerges gradually from the machine operations and learning related to its goals and values. These are built using a self-organizing hierarchy of spatio-temporal patterns that result from learning useful actions.

Neurons in the motor pathway are organized in a semi-hierarchical way to be able to store a large number of skills and actions. At the bottom of the hierarchy are row motor outputs. Memories are developed through registering these activities by modifications of the interconnection weights between neurons on the lower levels and activated neurons on the higher level.

As in the sensory pathway, hierarchies of action representations are built bottom-up, from the simplest actions that require little sensory-motor coordination or sequential memories, to the most complex ones, that may last for a long period of time and require lots of memories and sensory-motor coordination. An example of a simple action may be a reactive response to a painful shock, while driving home may be an example of a more complex action. Higher level actions can only be obtained after lower level skills are learned.

Similar reasons that were presented to the sensory pathway can be given to justify a hierarchical organization of motor actions. First is the large capacity to learn various skills. In the motor pathway hierarchy, the number of motor neurons on the higher levels is much greater than that of the lower levels; however, the number of activated neurons that represent skills and actions on the higher levels is less than that on the lower levels. Lower levels may no longer be capable of storing any new information since they were involved in learning many patterns. Instead they represent simpler actions that are used to implement actions on the higher levels. This results in lower plasticity of the interconnections on the lower level and higher plasticity on the higher levels.

EI activates motor neurons in response to a request for action from the value system in coordination with the sensory neuron activity representing the perceived state of the EI. If an action was taken that resulted in a positive value, an association between the sensory and the motor neuron's activity is learned. This makes it more likely that a similar action will be executed again when the EI is exposed to a similar environmental situation. The same complex operation may be executed using various simpler operations, which leads to a similar concept of invariance building in the motor pathway as that in the sensory pathway. Continuity of a higher level action is used for invariance building between an action represented on the higher level and its lower level implementations.

A prediction mechanism is also used in the motor pathway. At every step of a motor action, a prediction is made regarding expected inputs from the sensory pathway and the value system. If the prediction is correct there is no need to learn any new associations.

Fig. 4 shows a schematic representation of interactions between sensory and motor units on different levels of HSOLM hierarchy. In this figure, representation connections indicate entity recognition along the upward sensory pathway and represent downwards activations along the motor pathways. Feedback connections represent expectations of future inputs in the sensory and motor pathways. These expectations are provided by both motor neurons as well as higher level sensory neurons. The direction arrows indicate stimulation links from sensory and goal creation pathways to motor neurons, while planning arrows connecting motor neurons to the sensory neurons predict the sensory inputs after the action represented by this motor neuron was completed.

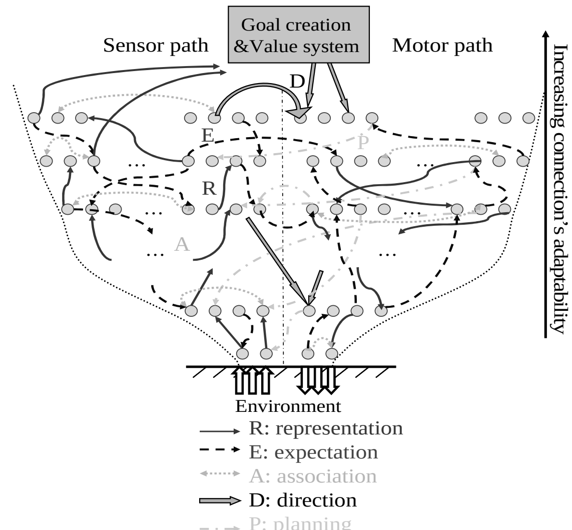

Fig. 4. Sensory-motor coordination links in HSOLM

3)Goal creation pathway

The primary objective of neurons in the goal creation pathway is creating goals, evaluation of actions in relation to its goals, learning the useful associations and stimulating the machine to perform useful actions. Goals are created based on external signals and internally generated stimuli that prompt the machine to do a desired action. A positive outcome of these desired actions will be the reason to learn internal representations, skills and useful associations.

Similar to neurons in the sensory and motor pathways, neurons in the goal creation pathway are organized hierarchically in order to represent a large number of goals and the means of their realization. Lower level goals are externally stimulated through special type of sensory inputs. Neurons' activation on these inputs may represent a large number of situations that the EI encounters while interacting with the environment. Higher level goals are developed through associations between activities on these lower level goal creation neurons and other neurons in the sensory-motor pathways. Hierarchical organization helps to obtain high goal capacity that the machine can develop through its interaction with the environment.

The goals emerge gradually from the machine operations. By using its goal creation system the machine defines higher level goals and determines the ways to implement them. The goal creation system generalizes the information received, introducing some degree of invariance for its higher level goals which increases with the level of generalization.

In the goal creation hierarchy, goals represented on the lower levels correspond to simple, externally driven objectives, while those on the higher levels correspond to complex objectives that are learned over the machine's lifetime and are related to its understanding of the best ways to accomplish the lower level goals. Since the lower level goals may be satisfied in many different ways, the correlations that result from their satisfaction are well defined and therefore less variable. We observe a similar level of increasing plasticity of the goal driven correlations on higher levels as we observed in the sensory and motor pathways. The machine self-organizes its goal creation process to formulate most of its goals on the higher levels of hierarchy.

As in the sensory pathway, the EI uses two mechanisms to create higher level goals and build invariances to implement these goals. One is screening for novelty. Only new and useful goals are stored. To do so the machine continuously predicts the level of satisfaction of the lower level goals. This prediction manifests its expectations regarding the lower level goals and an assumption of sameness of the rules under which such goals can be satisfied.

When EI is situated in a specific environment, its goals can be satisfied in a certain way, and the rules for achieving its objectives do not change drastically. Stationarity of the rules that affect goal satisfaction are used for invariance building of higher level goals and the way they are implemented (through various **subgoals**). For instance if we do not want to be hungry, we should keep our job to get money for food. Keeping the job is this higher level objective related to not feeling hungry. It must not be easily changed by temporary signals of feeling hungry that can be satisfied by lower level goals (like buying and cooking the food).

Goal creation is less understood than the other two pathways of EI. Therefore, in the following sections we devote more attention to developing the concept and structures for the goal creation pathway.

C. Goal Creation System (GCS)

1) Fundamental characteristics of the goal creation system

Before presenting details of the proposed goal creation system (GCS), we will discuss the basic desired features of the GCS in intelligent machines. First, an agent must be forced to explore the environment seeking solutions to achieve its goals. By interacting with the environment the agent accumulates the knowledge about the environment and its own embodiment. Without interaction with its environment, the machine will not be able to accumulate knowledge, develop its skills, nor formulate and implement its goals. Without such interaction we would not be able to call the machine intelligent.

Second, we assume that EI value system will evolve through self-organization, association, and learning using a simple built-in mechanism. We propose to base such mechanism on dedicated sensory inputs, called **“primitive pains”**, which trigger learning of values. The “primitive pain” inputs include pain, hunger, urge, discomfort and other special signals. The agent has a desire to reduce the pain or equivalently to pursue pleasure/comfort. Pain detection and moderation of pain by detecting a change in the pain (pleasure) intensity will represent the most primitive values. We imply that such requirements will make the proposed goal creation system biologically plausible even at the level of human intelligence. In a new-born baby, a hierarchical goal creation system and value system has not been yet developed. If he is exposed to a primitive pain and suffers, he will not be satisfied until the pain signal is reduced. When the pain is reduced the baby learns objects and actions that helped to lower that pain.

Since the pain signal comes from the environment (including embodiment of EI), we must make sure that it is inevitable unless the machine learns how to avoid it. Thus a mechanism must be put in place to gradually increase the pain level unless the machine does something to reduce its pain. This is needed to prevent the machine from choosing not to do anything or to quit trying. Thus, an unstable environment perception (in this case in the form of the increasing pain level) is a necessity for learning in this model.

Pain reduction is desirable while pain increase is not. Since the machine has several primitive pains, each one of them has its own changing intensity, and requires its own solution. At any given time, the machine suffers from the combination of different pains of different intensity, as shown in Fig. 5.

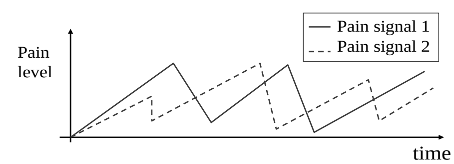

Fig. 5. Changes in temporal intensity of primitive pain signals.

It is easy to make references to biological systems where a similar mechanism is used to induce activity-based exploration and learning. For instance, one of the fundamental instincts we have in order to survive is to find and eat food to sustain our activities. A gradually increasing discomfort coming from our sensors in the stomach tell us that we must eat. Similar urges pressure us to go to the bathroom, use clothes when it is cold, or not touch a burning coal. The pain warns us against incoming threats, but also forces us to take an action. We also feel relief if we take an action that reduces this pain. Thus pleasure can be perceived as opposite to pain.

The intensity of the perceived pain may serve as a regulator to set priorities on our actions and thus be responsible for goal creation. For example, the urgent need to go to the bathroom may easily overtake our desire to eat, or even more so to sit through an interesting lecture. In general, the strongest pains will determine the most pressing goals. Thus the pain based GCS will also yield a goal management scheme.

A primitive pain may be responsible for primitive goals, but it can also be a trigger for developing higher level pain/pleasure centers and a mechanism for the creation of higher level goals. This is based on a fundamental mechanism for need to act and a simple measure for satisfying such a need. We would like to argue that this simple need to act may lead to complex goal creation and its implementation. The mechanism of goal creation in a human is not obvious to behavioral scientists or to psychologists. It is likely that the mechanism we propose has nothing to do with the way people create their goals. However, it is biologically feasible, simple, and it satisfies our need to establish goal creation for machine learning. In addition, this goal creation system stimulates the machine to act.

2)Basic unit of GCS

The proposed goal creation mechanism is a simple one, yet it may lead to formulation of complex goals and checks for their implementation. This mechanism is based on three groups of neurons that interact with each other and with the machine's memory. In the first group, the **pain detection center** (showing the pain level) is stimulated by the sensory inputs and represents the negative stimulation, pain, discomfort, or displeasure. A **delayed pain center** stores the delayed pain level. The second group of neurons compares signals from the pain detection center and the dual pain center and registers a decrease or increase in the pain level. This group sends a positive or a negative reinforcement to learn the sensory-motor coordination. Finally, the third group contains active neurons in the sensory and motor pathways of the EI.

Outputs of the second group of neurons correspond to the reinforcement learning signals used to instruct the machine regarding desired or undesired state/action pairs. In a classical reinforcement learning control structure, state/action pairs are evaluated by a critic network and the machine adjusts its actions to optimize the output value of the critic network. A critic network learns the values of state-action pairs through the reinforcement signals it receives from the environment or a teacher. While, in the proposed GCS, neurons in the second group, which monitor changes in the activity of the pain levels, create the reinforcement signals automatically. Their activation indicates either a positive or a negative value. The third group are normal sensory and motor neurons. The basic goal creation cell structure is shown in Fig. 6. This cell triggers exploration for a proper action and learning.

A gradually increasing pain level forces the machine to explore, since this is the only chance that the machine will learn a proper action when it has yet to learn anything about its environment. The machine explored its environment by switching among motor actions. This exploration comes with no learning until a reward (positive or negative) is received. Once the pain reduction or increase is detected by the second group of neurons, a learning signal is produced to reinforce or weaken the value of an action by strengthening or weakening its interconnection links.

Pain increase will be a control signal for currently activated motor signals not to fire (thus making the sensory-motor link more inhibitory). Pain reduction will make links between active sensory and motor neurons more excitatory. The stronger the change in the pain level, the stronger the reinforcement signal and the link weight adjustment is. Therefore, the neurons in the second group act as reinforcement neuro-transmitters stimulating the machine learning. In this model we assume that a positive reinforcement satisfies one of the goals, and is perceived as a reduction of a pain signal.

This mechanism alone is not enough. After several unsuccessful motor actions, all sensory-motor links would become inhibitory and the machine would quit doing anything. To prevent this from happening, a direct link from the pain center and the increasing level of pain force exploration. The described interaction of various groups of neurons in the goal creation mechanism is illustrated in Fig. 6.

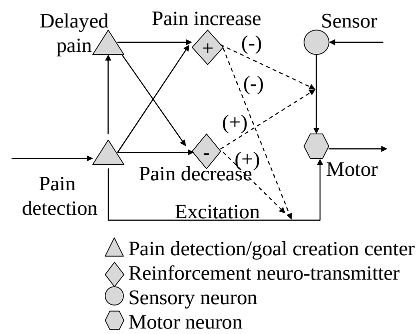

Fig. 6. Basic goal creation cell and learning of sensory-motor coordination.

This simple mechanism is easy to expand and generalize. In order to generate abstract and complex goals, we will incorporate basic goal creation units into different hierarchical levels of the goal creation pathway as discussed next.

3)Building goal hierarchy

A primitive pain is a signal received from the primitive pain sensors. It stimulates the primitive pain detection center. An agent is thus stimulated to explore for actions or to exploit the action that relieves its pain. The exploration is based initially on the random associative links or links that were initially (genetically) set to facilitate the reduction of the primitive pains. Such genetically set links facilitate learning of higher level skills and correspond to built-in skills. This may be a preferred solution to designing machines that need to develop complex skills.

Genetically set associations between the primitive pain centers and actions also exist in humans. A baby cries when it is wet or hungry, it also has well developed sucking reflexes to eat. A burning pain from touching a hot plate triggers an automatic pull back reflex. These sensations and actions become gradually associated with circumstances under which they occurred, leading an intelligent agent to learn basic skills or improve upon them.

By detecting the change in the pain signal level reinforcement neuro-transmitters are activated and send out the positive or negative reinforcement signals to change the weights of the sensory-motor connections. After several random trials, the action "eat", connected with perception of "food", will be rewarded. As a result, the strength of links from "food" and "hunger" to "eat" will be increased. Whenever the "hunger" pain center sends out pain signals, the "eat" will be excited prompting machine for this action.

Primitive pain centers are connected through excitatory links to a number of **abstract pain centers**. Each time the primitive pain center is excited, it sends activation signals to the abstract pain centers linked to this primitive pain center. Thus abstract pain centers echo the primitive pain. However, since these centers are not stimulated from the original pain sensors, they only symbolize the real pain. They can be inhibited by sensory neurons that are associated with elimination of the primitive pain.

When "food" is available and the agent "eats", this suppresses the primitive pain as well as the abstract pain. The pain signal disappears and the agent goes back to its normal state. As a result an inhibitory link is developed between sensory signal "food" and the abstract pain center. When "food" is not available, the agent cannot reduce the real pain. However, he may be able to do something to reduce an "abstract pain". Thus a new, higher level goal is created to reduce the "abstract pain". Although reduction of the abstract pain (getting "food") does not directly reduce the real pain ("hunger") it is may be a prerequisite for such reduction.

The agent is forced to explore to solve the abstract pain. Again, exploration is done based on initial associative connections between sensory and motor pathways. The abstract pain center will force this exploration. The reinforcement transmitters connected with abstract pain center change the interconnection weights. Eventually, the reduction in the abstract pain that indicates no "food" will be associated with the sensory-motor pair "refrigerator"-"open". It does not matter if such action (opening refrigerator) was found by pure exploration or by instruction from a teacher, as long as it is undertaken, it will be reinforced. Of course, by exploration and reinforcement, the action "open" will be associated with the perception "refrigerator" and the abstract pain signal "no food", since once the agent opens the refrigerator, he sees the food and the abstract pain is suppressed. In addition, since there is an association between "refrigerator-open" and "food", "food" will be expected through expectation link from the sensory input "refrigerator" and motor unit "open" to the sensory unit "food". This expectation link will be used for planning future actions in which reducing this kind of pain may be required. Such process can be illustrated using Fig. 7.

Normally, the goal creation process is much more complex with many goals created. The goals on different abstraction levels form the hierarchy of goals in the goal creation system. Following the previous example, if the agent "opens" the "refrigerator", but the "food" is not found, the agent needs other options to suppress the abstract pain, and subsequently the primitive pain. He may explore by random search or by instruction. Once he "spends" some "money", to buy the food the abstract pain (no food) is reduced; such action is rewarded and more strongly connected with abstract pain center.

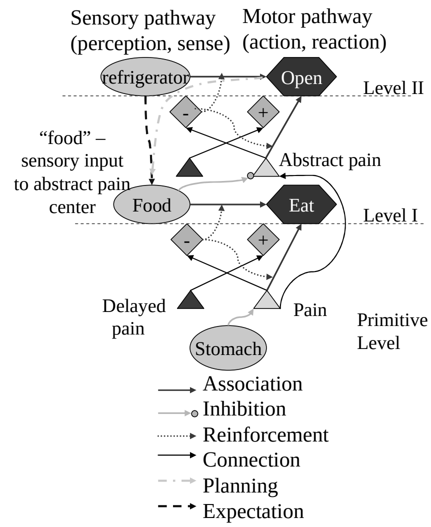

Fig. 7. Reduction of abstract pain signal

When "money" is available and spent, the "food" is obtained. The "food" is eaten, the primitive pain is suppressed, and the pain signals are reduced. Otherwise, when "money" is not available, an abstract pain center on level II is activated with an inhibitory link from "money". Subsequently, when "money" is not available, the agent needs to do something to solve this problem and reduce pain on level II. The pain center on level II is directly stimulated by the pain center on level I. Again, a pain center on level II represents a memory of pain on level I and has automatically but slowly increasing excitation level. After exploration, the agent finds out that the solution to the pain represented by "no money" is to "work" at a "job". The agent may also find that "stealing" other person's "purse" can provide "money". However, even if such action suppress the pain on this pain branch, the agent will be punished by inflicting pain on other pain branch so that this association will be weakened and "working" at a "job" will stand out as the best option. Subsequently, pain centers on higher levels will be created and the hierarchy of pain centers and goal creations will be built.

Accordingly, instead of a computational-based value system used in typical reinforcement learning, the value system is essentially embedded in such goal creation system. Initially, the agent acts on very primitive goals and learns the sub-goals or higher level goals through interaction with its environment. At every step, the agent finds an action that satisfies its goals and such action may result in creating further goals. Gradually the agent learns values of various states for implementing goals. It also learns to associate the primitive goals with its internal states and learns to create higher level goals. At more advanced levels, the agent is able to understand external instructions and use them as its goals.

D.Goal-driven learning system

Based on the previous discussion, it is noticed that the neurons in GCS form the pain centers on different abstraction levels. These pain centers produce the excitation to the motor signals on different levels. Therefore, in the proposed learning system, sensory pathway, motor pathway, and the goal creation pathway all have the hierarchical structures that interact with each other. In addition, it is stated that the machine may be stimulated from different primitive pains. Each primitive pain has its own hierarchical pain tree. It is possible that different pain trees overlap on certain pain branches. The motor neurons will accept stimulations from different pain trees and the strongest pain signals influence the choice of action. The motor pathway will respond to the pain signals from the goal creation pathway and to specific situation represented by the sensory pathway to initiate the desired action. The structure of goal-driven learning system is schematically illustrated in Fig. 8.

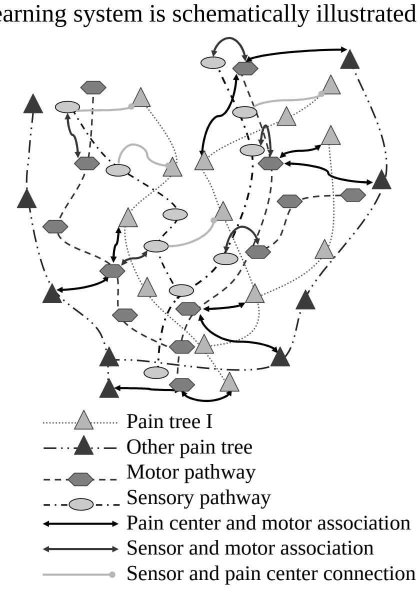

Fig. 8. Goal-driven learning scheme

During the process of goal creation, the reinforcement signals play a significant role in building sensory-motor associations. The reinforcement signals come from the environment, and include autonomous primitive pain signals as well as the teacher's input. The agent's experience built through interaction with the environment affects the associations between different sensory inputs and different goals. The tree pathways develop simultaneously. First, low level recognition in the sensory pathway and simple motor actions are developed to manage the primitive pain signals. The machine focuses its attention on objects and actions it can use to lower the pain; thus learning is selective. Not all the objects in the sensory inputs are of interest to the machine, and only those that are will be represented in its memory.

Once representations that are associated with lower level goals are formed and the machine is capable to formulate higher level goals, it may extend its representations to include higher level concepts on the sensory pathway, and learn a sequence of actions to implement them. Thus memories of more abstract entities and useful actions will be formed on the higher levels. Gradually, the complexity of the machine's operations - its intentional memory of objects, actions and goals increases; and its three hierarchical pathways are being built. As a result, complex, long lasting goals may be created and managed by such system, resulting in a complex, intelligent behavior.

IV. Hardware vs. Software for EI

Embodied Intelligence can be implemented in hardware or software. In the following comparison, hardware implementation will use multiple processors working concurrently, while software will run on a single CPU. This is not to say that each processor in hardware implementation will represent a single neuron. The processing speed of today's hardware is much higher than needed to simulate a single neuron in a real time. Thus in hardware implementation a hybrid approach can be used, where a group of neurons is simulated at each concurrent processor, and a number of neurons simulated is set to deliver a real time operation. This depends on the speed of operation. For instance, if a concurrent hardware operates at 200 MHz, a single operation may take 5 nsec, so in 5 msec needed for real time neural response, and then up to 1mln operations can be performed. This will set a limit on the number of neurons that such concurrent hardware can simulate.

Software implementation is a convenient choice today due largely to inadequate hardware design or programming tools for easy implementation and experimentation with cognitive mechanisms. However, software simulation has inherent limitations for implementing real-time operation of EI with brain level complexity. The major limitation comes from replacing the network of interconnected concurrent processors by a single central processing unit (CPU). Not only does the CPU have to run n times faster to compensate for the combined processing power of n neurons, but also it must simulate the complexity of the interconnections. With the average number of interconnections growing with the size of the network and the time the machine spends updating the interconnections and simulating signal transformations through these interconnections dominate. For instance, to simulate the human brain with 10^{11} neurons and each with an average of 10000 connections per neuron, the CPU must run 10^{15} times faster than biological neurons. With the average response time of a neuron on the level of 5msec, the CPU would have to perform 10^{15} operations in 5 msec. Assuming that a single operation can be performed in one clock cycle, this would require the clock speed of 200,000,000 GHz, (or 10mln times faster than current computers). During the time period that corresponds to such switching frequency light travels on the distance of 1.5 nm. Thus even if the switching is performed with the speed of light, the device geometry should be comparable with the particle size of silicon. Operating on such small scale would require going beyond the single electron switching or spin electronics. Thus a single CPU may never be able to perform the real-time operations of system with human brain complexity. It is very likely that its use would have to be limited to modeling small sized networks necessary to develop mechanisms for EI. It is our opinion that only hybrid hardware approach may reach complexity of human level intelligence in the predictable future.

There are other reasons why hybrid hardware implementation of EI is preferable over software simulation on a single CPU. We will just mention two major ones.

First, the software operation is synchronized by a central clock. To perform a single interconnection weight update or to do a single neuron operation the clock that is powering all the circuits is dissipating huge amounts of energy. The essential computing is performed by changing a signal value on a small load capacitance, and the related switching energy that must be delivered to do so is almost negligible comparing to the clock energy loss. Even if some power savings techniques are used, the average energy loss per operation in the CPU is large. Currently computers running on 3 GHz clock dissipate over 100 W of power. Increasing the clock frequency will significantly increase this energy loss. By comparison, using multiple concurrent processors to perform the same number of operations as simulated in software would require much lower frequency, therefore would consume much less power.

Second, the hardware implementation of EI with millions of processors working concurrently provides robustness against hardware failures, similar to the one observed in the human brain. If a single processing unit (out of millions) is damaged, it will have only a minor effect on the overall performance of the EI. Neighboring processors will take over its function in a similar way that exists in the human brain. This would allow for using wafer scale integration to deliver hardware devices for such concurrent systems, reducing the manufacturing and testing cost. In fact, many processing units in the wafer scale array of processors may malfunction, and the system may still operate equally well, providing that there are still many operational processors in the system. On the other hand, a software system is very sensitive to hardware errors, and a single fault can make such system not operational.

A major advantage of the software approach is that it is relatively low cost and uses well developed programming methods. Hardware prototypes are harder to built and carry a significant development cost. For this reasons software approach is very popular and may be used for many years to study and develop models of intelligence.

V. Conclusions

In this paper we presented a framework to design working models for embodied intelligence. Structural elements of this framework came from our definition of intelligence and its embodiment. Three self-organizing hierarchical structures – sensory, motor, and goal creation pathways were used to form the core of EI. The three pathways are developed simultaneously extracting knowledge and learning skills through interaction with the environment. These three pathways interact on various abstraction levels, providing associations among them.

The goal creation pathway is responsible for goal creation, evaluation of actions in relation to its goals, learning of useful associations, and stimulating machine to perform useful actions. It also stimulates the growth of hierarchical structures representing sensory inputs, actions and skills acquired by the machine, and abstract goals that machine creates for itself. EI learns predominantly in unsupervised manner by responding to stimuli from the environment. Learning is deliberate, perpetual, and related to satisfactory completion of EI goals. Hostile stimulation from environment is necessary for EI to grow in sophistication and to acquire necessary knowledge and skills.

Technological challenges of building human-level intelligence are briefly discussed by comparing properties of software and hardware implementations. However, it is our strong feeling, that technology which supports hybrid implementation of EI must be developed and used to design human level intelligence. Such technology will benefit society and stimulate economical growth for the years to come.

References

[1] R. A. Brooks, “Intelligence without reason,” *Proc. IJCAI-91*, pp. 569-595, 1991.

[2] R. Pfeifer and C. Scheier, *Understanding Intelligence*, MIT Press, Cambridge, MA, 1999.

[3] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, MA, 1998.

[4] R. A. Brooks, *Flesh and Machines: How Robots Will Change Us*, Pantheon, 2002.

[5] O. Sporns, and G. M. Edelman, “Solving Bernstein's problem: A proposal for the development of coordinated movement by selection.” *Child Development*, *64*(4), 1993, pp. 960-981.

[6] J. Hawkins and S. Blakeslee, *On Intelligence*, Times Books, 2004.

[7] S. Card, T. Moran, and A. Newell, *The psychology of human-computer interaction.* Hillsdale, NJ: Lawrence Erlbaum Associates, 1983.

[8] A. Newell, *Unified Theories of Cognition*. Cambridge, Massachusetts: Harvard University Press, 1990.

[9] S. Nason and J. E. Laird, “Soar-RL, Integrating Reinforcement Learning with Soar,” *Int. Conf. Cognitive Modeling*, 2004.

[10] J. Anderson, *Rules of the Mind*, Lawrence Erlbaum Associates, Hillsdale, NJ, 1993.

[11] V. B. Mountcastle, “Introduction to a special issue of cerebral cortext on columns”, *Cerebral Cortex*, 13, pp. 2-4, 2003.

[12] J. DeFelipe, and E. G. Jones, *Cajal on the Cerebral Cortex.* An Annotated Translation of the Complete Writings, Oxford Univ. Press, 1998.

[13] E. G. Jones, Microcolumns in the Cerebral Cortex, *Proc. National Academy of Science of USA*, vol. 97(10), pp. 5019-5021, 2000.